

Classification and clustering for samples of event time data using non-homogeneous Poisson process models

Duncan S Barrack^{a*} and Simon Preston^b

^a *Horizon Digital Economy Research Institute, University of Nottingham, Nottingham, UK.*

^b *School of Mathematical Sciences, University of Nottingham, Nottingham, UK.*

Abstract

Data of the form of event times arise in various applications. A simple model for such data is a non-homogeneous Poisson process (NHPP) which is specified by a rate function that depends on time. We consider the problem of having access to multiple independent samples of event time data, observed on a common interval, from which we wish to classify or cluster the samples according to their rate functions. Each rate function is unknown but assumed to belong to a small set of rate functions defining distinct classes. We model the rate functions using a spline basis expansion, the coefficients of which need to be estimated from data. The classification approach consists of using training data for which the class membership is known and to calculate maximum likelihood estimates of the coefficients for each group, then assigning test samples to a class by a maximum likelihood criterion. For clustering, by analogy to the Gaussian mixture model approach for Euclidean data, we consider a mixture of NHPP models and use the expectation maximisation algorithm to estimate the coefficients of the rate functions for the component models and probability of membership for each sample to each model. The classification and clustering approaches perform well on both synthetic and real-world data sets considered. Code associated with this paper is available at <https://github.com/duncan-barrack/NHPP>.

Keywords: point process; non-homogeneous Poisson process; classification; clustering; expectation maximisation.

1 Introduction

Much real world data is amenable to modelling as a Poisson process (Ross et al., 1996), for example the times at which a customer makes purchases in a shop (Schmittlein et al., 1987), times of requests to internet web servers (Arlitt and Williamson, 1997) and the firing times of neurons (Brown et al., 2002). In a wide range of applications one observes multiple samples of such data on a common interval and a challenge is to label each sample as belonging to one of a small number of distinct classes. For example, for the purposes of targeted marketing, marketers may wish to label customers according to their temporal shopping habits. Furthermore, by categorising stores by till transaction times retailers can gain valuable insight in to the different demands placed on different store types. This in turn can help to inform various

*Corresponding author at: Horizon Digital Economy Research Institute, University of Nottingham, Nottingham, NG7 2TU, UK. Tel.: +44 (0)115 8232554. E-mail address: duncan.barrack@googlegmail.com.

decisions from stock management to appropriate staffing levels to meet demand which varies between different types of stores and across the day. Similarly transportation companies can benefit from the ability to categorise transportation hubs such as railway stations according to the times at which customers enter to help inform the kind and number of transportation services they need to offer at each type of hub. When a ‘training set’, which contains samples with known class memberships, is available such a task is termed ‘classification’, when no labels are available for the data it is termed a ‘clustering’ problem (Hastie et al., 2009).

Previous approaches to classification and clustering of temporal data of this type have focussed on deriving summary statistics from the data such as the mean and variance of the inter-event times and using these as inputs in statistical learning models (Moore and Zuev, 2005; Roughan et al., 2004; McGregor et al., 2004). In addition to studies of temporal data a number of authors have considered spatial point processes. Illian *et al* used functional principal component analysis on second order summary statistics of spatial point process data to classify plant species (Illian et al., 2006, 2004). Closer in approach to the present work Linnett et al. (1995) used a maximum likelihood (ML) criterion to classify the textures of images modelled using homogeneous Poisson processes (HPPs).

HPPs, like those used in the work of Linnett et al. (1995) assume rates are constant for each class. However, there are numerous empirical examples of event data which can not be adequately modelled by HPPs (e.g. stock transaction times and human communication behaviour both been shown to exhibit periodic patterns (Engle, 2000; Malmgren et al., 2008), computer network traffic are better modelled by NHPPs rather than HPPs (Paxson and Floyd, 1995)), and this assumption is too restrictive for many types of data which crop up in the real world. Hence, in this work, we consider instead non-homogeneous Poisson process (NHPP) models, which are defined in terms of rate functions that vary with time. We model these rate functions using a spline basis expansion which imposes an assumption of smoothness on the functions and means estimating them amounts to estimating a finite number of basis coefficients. We go beyond previous work which has focussed purely on rate function estimation (Alizadeh et al., 2008; Kuhl and Bhairgond, 2000; Scott et al., 1980; Zhao and Xie, 1996; Massey et al., 1996; Lee et al., 1991). In particular, after proving that the parameters of the rate function representations can be found via maximum likelihood estimation by solving a convex optimisation problem, we consider the tasks of classification and clustering. For classification we show that once intensity estimates for the NHPP rate functions corresponding to each class are obtained, it is a simple task to obtain the posterior class membership probabilities for test data. For clustering, where class labels are not known *a priori*, we assume all data to have been generated from a mixture of k NHPP models. Again the rate functions for each model are represented by a combination of basis functions but this time the expectation maximisation (EM) algorithm is used to obtain the coefficients of the basis functions as well as k membership probabilities for each distinct sample of event times. Our approaches are validated on synthetic and two real data sets.

The paper is structured as follows. We begin by describing our procedure for estimating NHPP rate functions in Section 2. In Sections 3 and 4 we outline our classification and clustering methods. The results of applying these to synthetic and real data are shown in Sections 5. We conclude with a discussion in Section 6.

2 Rate function estimation

Suppose that we assume that a sample of event times $\chi = \{\chi_1, \chi_2, \dots, \chi_{m_\chi}\}$, observed on the interval $\{t : t_1 < t \leq t_2\}$, arises from a NHPP with rate function $\lambda(t)$, then the log likelihood for the rate function is

$$l(\lambda(t)|\chi) = -M(t) + \sum_{j=1}^{m_\chi} \ln(\lambda(\chi_j)), \quad (1)$$

where $M(t) = \int_{t_1}^t \lambda(\tilde{t}) d\tilde{t}$ (see Thompson (2012)).

In finding our maximum likelihood estimate for $\lambda(t)$ in equation (1), we impose that the rate function belongs to a class of smooth functions by assuming $\lambda(t)$ is a linear combination of smooth basis functions

$$\lambda(t) = \sum_{m=1}^{n_b} c_m B_m(t). \quad (2)$$

Here $B_1(t), \dots, B_{n_b}(t)$ could be any basis function. However, as we only consider event times defined over a finite interval in this work, henceforth we use a spline basis (De Boor, 1978) which has finite support (as opposed to, for example, a Gaussian function which has infinite support) and is therefore a natural choice to model NHPP rate functions defined over a fixed interval. We estimate the basis function coefficients $\mathbf{c} = \{c_1, c_2, \dots, c_{n_b}\}$ by minimising $-l(\lambda(t)|\chi)$ i.e. by minimising

$$f(\mathbf{c}) = -l(\lambda(t)|\chi) = M(t) - \sum_{j=1}^{m_\chi} \ln \left(\sum_{m=1}^{n_b} c_m B_m(\chi_j) \right), \quad (3)$$

subject to the constraint that $-\lambda(t)$ is less than or equal to zero, i.e.

$$g(\mathbf{c}) = -\lambda(t) = - \sum_{m=1}^{n_b} c_m B_m(t) \leq 0. \quad (4)$$

This ensures that the rate function estimate $\hat{\lambda}(t) = \sum_{m=1}^{n_b} \hat{c}_m B_m(t) \geq 0 \forall t$ which is consistent with the requirement that NHPP rate functions are non-zero for all values of time.

By showing that $f(\mathbf{c})$ and $g(\mathbf{c})$ are twice differentiable and convex in \mathbf{c} , the minimisation to perform is straightforward using a technique suitable for non-linear convex optimisation problems with an inequality constraint such as the interior point method (Byrd et al., 1999, 2000; Waltz et al., 2006). We show this by first proving (3) is convex in \mathbf{c} . The gradient of equation $f(\mathbf{c})$ is

$$\nabla f = (\nabla_1(\mathbf{c}), \dots, \nabla_{n_b}(\mathbf{c})), \quad (5)$$

where $\nabla_i(\mathbf{c}) = \frac{\partial}{\partial c_i} M(t) - \sum_{j=1}^{m_\chi} \frac{B_i(\chi_j)}{\sum_{m=1}^{n_b} c_m B_m(\chi_j)}$. The $n_b \times n_b$ Hessian is

$$\mathbf{H} = \begin{bmatrix} H_{1,1}(\mathbf{c}) & \dots & H_{1,n_b}(\mathbf{c}) \\ \vdots & \ddots & \vdots \\ H_{n_b,1}(\mathbf{c}) & \dots & H_{n_b,n_b}(\mathbf{c}) \end{bmatrix}, \quad (6)$$

where $H_{i,u}(\mathbf{c}) = \sum_{j=1}^{m_\chi} \frac{B_i(\chi_j)B_u(\chi_j)}{(\sum_{m=1}^{n_b} c_m B_m(\chi_j))^2}$ as $\frac{\partial^2}{\partial c_i \partial c_u} M(t) = \int_{t_1}^t \left(\frac{\partial^2}{\partial c_i \partial c_u} \sum_{m=1}^{n_b} c_m B_m(\tilde{t}) \right) d\tilde{t} = 0$ for every i and u .

As we may express $\mathbf{H} = \mathbf{X}'\mathbf{X}$, where

$$\text{the } \chi_{m_\chi} \times n_b \text{ matrix } \mathbf{X} = \begin{bmatrix} \frac{B_1(\chi_1)}{\sum_{m=1}^{n_b} c_m B_m(\chi_1)} & \dots & \frac{B_{n_b}(\chi_1)}{\sum_{m=1}^{n_b} c_m B_m(\chi_1)} \\ \vdots & \ddots & \vdots \\ \frac{B_1(\chi_{m_\chi})}{\sum_{m=1}^{n_b} c_m B_m(\chi_{m_\chi})} & \dots & \frac{B_{n_b}(\chi_{m_\chi})}{\sum_{m=1}^{n_b} c_m B_m(\chi_{m_\chi})} \end{bmatrix}, \quad (7)$$

\mathbf{H} is positive semi-definitive as $\mathbf{y}'\mathbf{H}\mathbf{y} = \mathbf{y}'\mathbf{X}'\mathbf{X}\mathbf{y} = (\mathbf{y}\mathbf{X})'(\mathbf{X}\mathbf{y}) = \|\mathbf{X}\mathbf{y}\|^2 \geq 0$ for every non-zero column vector $\mathbf{y} = \{y_1, y_2, \dots, y_{n_b}\}$ of real numbers. Therefore $f(\mathbf{c})$ (equation (3)) is convex in \mathbf{c} .

As $g(\mathbf{c})$ (equation (4)) is linear in \mathbf{c} it is both convex and concave in \mathbf{c} . Hence estimating the basis function coefficients of the NHPP rate function $\lambda(t)$ is a convex optimisation problem.

3 Classification

Consider a set of n_α training samples $\mathbf{T}_{\text{train}} = \{\alpha_1, \alpha_2, \dots, \alpha_{n_\alpha}\}$ where each sample $\alpha_i = \{a_{i,1}, a_{i,2}, \dots, a_{i,m_i}\}$ is comprised of m_i event times. Each sample α_i is assigned a class label $G_i \in 1, 2, \dots, k$. The classification task consists of using this training data to predict the class labels of a test set $\mathbf{T}_{\text{test}} = \{\beta_1, \beta_2, \dots, \beta_{n_\beta}\}$ where only the event times for each sample i.e. $\beta_i = \{b_{i,1}, b_{i,2}, \dots, b_{i,m_{\beta_i}}\}$ are observed. Assuming each training sample is independent, the log likelihood estimate for the rate function $\lambda_\nu(t)$ for class ν is

$$l(\lambda_\nu(t) | \mathbf{T}_{\text{train}}) = \sum_{l=1}^{n_\alpha} \mathbb{1}(G_l = \nu) \left[-M_\nu(t) + \sum_{j=1}^{m_l} \ln(\lambda_\nu(a_{l,j})) \right], \quad (8)$$

where the indicator function $\mathbb{1}(G_i = \nu) = 1$ if $G_i = \nu$ and 0 otherwise. The rate function $\lambda_\nu(t)$ is given by

$$\lambda_\nu(t) = \sum_{m=1}^{n_b} c_{\nu,m} B_m(t). \quad (9)$$

Estimates for the basis function coefficients of the expression are found via the maximum likelihood estimation outlined in Section 2.

Classification of test data

For each test sample β_i the posterior probability it is a member of class ν can be derived from Bayes' theorem as follows

$$P(G_i = \nu | \beta_i) = \frac{e^{-\widehat{M}_\nu(t)} \prod_{j=1}^{m\beta_i} \widehat{\lambda}_\nu(b_{i,j})}{\sum_{q=1}^k e^{-\widehat{M}_q(t)} \prod_{j=1}^{m\beta_i} \widehat{\lambda}_q(b_{i,j})}. \quad (10)$$

If observations are to be assigned to the class for which this probability is maximal, then this corresponds to the Bayes classifier which is optimal under 0-1 loss.

4 Clustering

For the unsupervised learning task of cluster analysis of event time data we assume the data to have been generated from a mixture of NHPP models. Furthermore, each sample is associated with a hidden latent variable which specifies which model, or 'cluster', it is from. Specifically, we assume n event time samples $\mathbf{D} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ are generated from a mixture of k NHPP models and let $\mathbf{z} = \{z_1, z_2, \dots, z_n\}$ be a set of latent random variables which determine the component from which each sample originates. For the mixture the complete data log likelihood function is

$$l(\boldsymbol{\theta} | \mathbf{D}, \mathbf{z}) = \sum_{l=1}^n \log(p(\mathbf{x}_l, \mathbf{z} | \boldsymbol{\theta})) = \sum_{l=1}^n \log \left(\sum_{q=1}^k \mathbb{1}(z_l = q) \tau_q \mathcal{L}(\lambda_q(t) | \mathbf{x}_l) \right), \quad (11)$$

and corresponding incomplete data log likelihood is

$$l(\boldsymbol{\theta} | \mathbf{D}) = \sum_{l=1}^n \log(p(\mathbf{x}_l | \boldsymbol{\theta})) = \sum_{l=1}^n \log \left(\sum_{q=1}^k \tau_q \mathcal{L}(\lambda_q(t) | \mathbf{x}_l) \right). \quad (12)$$

Here τ_ν is the 'mixing weight' of component ν , where $0 \leq \tau_\nu \leq 1$ and $\sum_{q=1}^k \tau_q = 1$. $\mathcal{L}(\lambda_\nu(t) | \mathbf{x}_i)$ is the likelihood for the rate function associated with NHPP model ν given the sample \mathbf{x}_i . As before, we represent the rate functions associated with each component ν as a linear combination of basis functions, i.e. $\lambda_\nu(t) = \sum_{m=1}^{n_b} c_{\nu,m} B_{\nu,m}(t)$.

We use the EM algorithm (see Dempster et al. (1977) for details) to find estimates for the parameters $\boldsymbol{\theta} = \{\boldsymbol{\tau}, \mathbf{C}\}$ (where $\boldsymbol{\tau} = \{\tau_1, \tau_2, \dots, \tau_k\}$, $\mathbf{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_k\}$ and $\mathbf{c}_\nu = \{c_{\nu,1}, \dots, c_{\nu,n_b}\}$ is the set of basis function coefficients for NHPP rate function ν) which maximise the log likelihood of the mixture. This procedure is described below.

Expectation step

Given a current estimate for the model parameters $\widehat{\boldsymbol{\theta}}^c$ we find the so called 'auxiliary function' which is the expected value of the data log likelihood (equation (11)) of the mixture model with respect to the conditional distribution of \mathbf{z} given the data \mathbf{D} .

$$Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^c) = \mathbb{E}_{\mathbf{z}|\mathbf{D}, \hat{\boldsymbol{\theta}}^c} [l(\boldsymbol{\theta}|\mathbf{D}, \mathbf{z})] \quad (13)$$

$$= \mathbb{E}_{\mathbf{z}|\mathbf{D}, \hat{\boldsymbol{\theta}}^c} \left[\sum_{l=1}^n l(\boldsymbol{\theta}|\mathbf{x}_l, z_l) \right], \quad (14)$$

$$= \sum_{l=1}^n \mathbb{E}_{\mathbf{z}|\mathbf{D}, \hat{\boldsymbol{\theta}}^c} [l(\boldsymbol{\theta}|\mathbf{x}_l, z_l)], \quad (15)$$

$$= \sum_{l=1}^n \sum_{q=1}^k p(z_l = q|\mathbf{x}_l, \hat{\boldsymbol{\theta}}^c) l(\boldsymbol{\theta}|\mathbf{x}_l, z_l), \quad (16)$$

$$= \sum_{l=1}^n \sum_{q=1}^k \hat{R}_{q,l}^c (\ln(\tau_q) + l(\lambda_q(t)|\mathbf{x}_l)). \quad (17)$$

Here $\hat{R}_{\nu,i}^c = p(z_i = \nu|\mathbf{x}_i, \hat{\boldsymbol{\theta}}^c)$ is the ‘membership probability’ of sample i to model, or ‘cluster’, ν . $l(\lambda_\nu(t)|\mathbf{x}_i) = -M_\nu(t) + \sum_{j=1}^{m_i} \lambda_\nu(x_{i,j})$ is the log likelihood for rate function ν for sample i . Plugging in the basis function expansions for the rate functions in to equation (17) we get

$$Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^c) = \sum_{l=1}^n \sum_{q=1}^k \hat{R}_{q,l}^c \left(\ln(\tau_q) - M_q(t) + \sum_{j=1}^{m_l} \ln \left(\sum_{m=1}^{n_b} c_{q,m} B_{q,m}(x_{l,j}) \right) \right). \quad (18)$$

To obtain the membership probabilities for of each sample i of model ν we regard the mixing weights $\boldsymbol{\tau}$ as prior probabilities of each mixture component and use Baye’s rule as follows

$$\hat{R}_{\nu,i}^c = p(z_i = \nu|\mathbf{x}_i, \hat{\boldsymbol{\theta}}^c) = \frac{p(z_i = \nu)p(\mathbf{x}_i|z_i = \nu, \hat{\boldsymbol{\theta}}^c)}{\sum_{q=1}^k p(z_i = q)p(\mathbf{x}_i|z_i = q, \hat{\boldsymbol{\theta}}^c)} = \frac{\hat{\tau}_\nu^c e^{-\hat{M}_\nu^c(t)} \prod_{j=1}^{m_i} \hat{\lambda}_\nu^c(x_{i,j})}{\sum_{q=1}^k \hat{\tau}_q^c e^{-\hat{M}_q^c(t)} \prod_{j=1}^{m_i} \hat{\lambda}_q^c(x_{i,j})}. \quad (19)$$

Maximisation step

After obtaining values for the membership probabilities, we obtain subsequent estimates for the parameters $\hat{\boldsymbol{\theta}}^{c+1}$ by finding the $\boldsymbol{\theta}$ which maximise the auxiliary function $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^c)$ (equation (18)). Note, we may maximise $\boldsymbol{\tau}$ and $\boldsymbol{\theta}$ separately, since they are not related. Updates for the mixing weights $\boldsymbol{\tau}$ are found by taking the partial derivative of equation (18), setting it to zero and solving for $\boldsymbol{\tau}$. The derivative, with respect to the mixing probability τ_ν of $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^c)$ is

$$\frac{\partial Q(\boldsymbol{\theta}, \boldsymbol{\theta}^c)}{\partial \tau_\nu} = \frac{\partial}{\partial \tau_\nu} \sum_{l=1}^n \sum_{q=1}^k \hat{R}_{q,l}^c \left(\ln(\tau_q) - M_q(t) + \sum_{j=1}^{m_l} \ln \left(\sum_{m=1}^{n_b} c_{q,m} B(x_{l,j}) \right) \right) + \delta \left(\sum_{q=1}^k \tau_q - 1 \right), \quad (20)$$

where δ is the Lagrange multiplier for the constraint that the mixing weights must sum to 1. Evaluating (20) and removing all terms not involving τ_ν we get

$$\frac{\partial Q(\boldsymbol{\theta}, \boldsymbol{\theta}^c)}{\partial \tau_\nu} = \sum_{l=1}^n \frac{d}{d\tau_\nu} \hat{R}_{\nu,l}^c \ln(\tau_\nu) + \frac{d}{d\tau_\nu} \delta \tau_\nu, \quad (21)$$

$$= \sum_{l=1}^n \frac{\hat{R}_{\nu,l}^c}{\tau_\nu} + \delta. \quad (22)$$

Setting equation (22) to 0 and solving for τ_ν we get

$$\tau_\nu = \sum_{l=1}^n \frac{\hat{R}_{\nu,l}^c}{-\delta}. \quad (23)$$

Using the constraint that $\sum_{q=1}^k \tau_q = 1$, $-\delta = \sum_{l=1}^n \sum_{q=1}^k \hat{R}_{q,l}^c = \sum_{l=1}^n 1 = n$. Hence the next estimate for the mixing weight $\hat{\tau}_\nu^{c+1} = \sum_{l=1}^n \frac{\hat{R}_{\nu,l}^c}{n}$.

Next we seek the \mathbf{C} which maximise $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^c)$ subject to the constraints that

$$\lambda_1(t), \dots, \lambda_k(t) = \sum_{m=1}^{n_b} c_{1,m} B_m(t), \dots, \sum_{m=1}^{n_b} c_{k,m} B_m(t) \geq 0. \quad (24)$$

This is a convex minimisation problem of the same variety as encountered in Section 2. To see this, consider the term $-f(\mathbf{c}_{\nu,i}) = -M_q(t) + \sum_{j=1}^{m_l} \ln(\sum_{m=1}^{n_b} c_{q,m} g(x_{i,j}))$ from equation (18). $f(\mathbf{c}_{\nu,i})$ is identical to equation (3) which we showed is convex in the set of basis function coefficients in Section 2. As $\hat{R}_{\nu,i} \geq 0$ for every ν and i , $\hat{R}_{\nu,i} f(\mathbf{c}_{\nu,i}) = \hat{R}_{\nu,i} \left(M_q(t) - \sum_{j=1}^{m_l} \ln(\sum_{m=1}^{n_b} c_{q,m} g(x_{i,j})) \right)$ is also convex. Furthermore, disregarding terms which do not involve \mathbf{C} , $-Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^c) = \sum_{l=1}^n \sum_{q=1}^k \hat{R}_{q,l}^c f(\mathbf{c}_{\nu,i})$ is convex in \mathbf{C} as sums of convex functions are also convex. The functions in the constraints given in equation (24) are linear in \mathbf{C} and are therefore convex and concave. Therefore to find the \mathbf{C} which maximise $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^c)$ (equation (18)) we can use the same numerical approach as we used to find the maximum likelihood basis function coefficients estimates in Section 2.

The expectation and maximisation steps are applied iteratively until the following condition is satisfied

$$Q(\boldsymbol{\theta}^f, \boldsymbol{\theta}^{f-1}) \leq Q(\boldsymbol{\theta}^{f-1}, \boldsymbol{\theta}^{f-2}) + \epsilon, \quad (25)$$

where $\boldsymbol{\theta}^f$ is the set of final parameter estimates and ϵ is a preset threshold.

Initialisation

Depending on the starting conditions the EM algorithm may converge to a local maximum. To overcome this issue we use a form of the random restart approach commonly used with the EM algorithm for Gaussian mixtures (Biernacki et al., 2003). For our procedure, to set initial estimates for \mathbf{C} , each sample \mathbf{x}_i is randomly assigned a class label $G'_i \in 1, 2, \dots, k$. Rate function estimates for the ‘random’ classes are obtained using the classification procedure outlined in Section 3 and

the basis function estimates which correspond to these are used for the initial estimates for \mathcal{C} . All mixing weights in τ are initially set to $1/k$. The EM algorithm is then initialised from a number of different sets of initial conditions corresponding to the different ‘random’ class assignments and the solution with the largest log likelihood is retained.

5 Empirical results

In this section we demonstrate the efficacy of our methods against synthetic and two real world data sets.

5.1 Synthetic data

We first apply our methods to synthetically generated event data with prescribed rate functions. None of the synthetic NHPP functions belong to the class of smooth functions we assume in our methods. We consider two separate two class problems and a four class problem which are described below.

Synthetic data set 1

Motivated by the many real world examples of point processes which exhibit periodic patterns (e.g. time of stock trades (Engle, 2000) and the times at which emails are sent (Malmgren et al., 2008)), for this set, sinusoidal rate functions are used for the NHPP rate functions. The frequency of oscillation for class 1 rate functions is half that of class 2, i.e.

$$\text{class 1 :- } \lambda_1(t) = 100\sin^2(t/2),$$

$$\text{class 2 :- } \lambda_2(t) = 100\sin^2(t).$$

Synthetic data set 2

Here, we mimic computer network traffic for which it has been shown that a NHPP model with a step function rate function is a suitable model (Paxson and Floyd, 1995).

$$\text{class 1 :- } \lambda_1(t) = \begin{cases} 20 & \text{if } t < \frac{t_2}{4} \\ 40 & \text{if } \frac{t_2}{4} \leq t < \frac{t_2}{2} \\ 60 & \text{if } \frac{t_2}{3} \leq t < \frac{3t_2}{4} \\ 80 & \text{if } t \geq \frac{3t_2}{4} \end{cases}$$

$$\text{class 2 :- } \lambda_2(t) = \begin{cases} 80 & \text{if } t < \frac{t_2}{4} \\ 60 & \text{if } \frac{t_2}{4} \leq t < \frac{t_2}{2} \\ 40 & \text{if } \frac{t_2}{3} \leq t < \frac{3t_2}{4} \\ 20 & \text{if } t \geq \frac{3t_2}{4} \end{cases}$$

Synthetic data set 3

Four classes of rate functions are considered by combining the rate functions of synthetic data sets 1 and 2.

$$\text{class 1 :- } \lambda_1(t) = 100\sin^2(t/2)$$

$$\text{class 2 :- } \lambda_2(t) = 100\sin^2(t)$$

$$\text{class 3 :- } \lambda_3(t) = \begin{cases} 20 & \text{if } t < \frac{t_2}{4} \\ 40 & \text{if } \frac{t_2}{4} \leq t < \frac{t_2}{2} \\ 60 & \text{if } \frac{t_2}{3} \leq t < \frac{3t_2}{4} \\ 80 & \text{if } t \geq \frac{3t_2}{4} \end{cases}$$

$$\text{class 4 :- } \lambda_4(t) = \begin{cases} 80 & \text{if } t < \frac{t_2}{4} \\ 60 & \text{if } \frac{t_2}{4} \leq t < \frac{t_2}{2} \\ 40 & \text{if } \frac{t_2}{3} \leq t < \frac{3t_2}{4} \\ 20 & \text{if } t \geq \frac{3t_2}{4} \end{cases}$$

5.1.1 Classification

For each class, 20 samples are generated using the thinning algorithm (Lewis and Shedler, 1979). 10 of these form the training set and the remainder form the out of sample test set. Using the classification procedure outlined in Section 3 we are able to recover the prescribed rate function for each data set (see Figures 1-3). Furthermore, when each sample in the test data is assigned to the class for which its membership probability (as defined in equation (10)) is maximal, the classification accuracy is 100%.

5.1.2 Clustering

Again, 20 samples for each class were generated. Using the clustering method detailed in Section 4 we were able to recover the NHPP rate functions for all data sets (see Figures 4 - 6). For each data set, we were able to correctly assign each sample to the NHPP model from which it was generated by assigning it to the model for which the membership probability (equation (19)) was maximal.

Synthetic data set 1 - classification task




Figure 1: Raster plots showing a selection of event times for each training and test sample for both classes together with plots illustrating that the rate functions of the two classes for synthetic data set 1 can be recovered using the procedure described in Section 3. $\lambda_1(t)$ (*resp.* $\lambda_2(t)$) is the prescribed rate functions of class 1 (2) and $\hat{\lambda}_1(t)$ ($\hat{\lambda}_2(t)$) is its estimate. 100 cubic basis functions (i.e. $n_b = 100$) were used for the rate function basis expansion given in equation (2) and all coefficients of these (i.e. $c_{1,1}, \dots, c_{1,100}, c_{2,1}, \dots, c_{2,100}$) were initially set to 1. The interior point method method was used for the optimisation (Byrd et al., 1999, 2000; Waltz et al., 2006).




Figure 2: Event times and rate function estimates for synthetic data set 2 illustrating that we are able to obtain accurate estimates for the NHPP rate functions for both classes.

Synthetic data set 3 - classification task




Figure 3: Selection of event times and rate function estimates for the final synthetic data set.

Synthetic data set 1 - clustering task




Figure 4: Raster plots of event times together with plots showing that the clustering method detailed in Section 4 can be used to recover the rate functions of the component models within a NHPP mixture for synthetic data set 1. A value of $\epsilon = 0.0001$ was used for the EM stopping criteria. Solutions from 10 different sets of initial conditions were obtained and the results for the solution with the largest likelihood were selected.

Synthetic data set 2 - clustering task

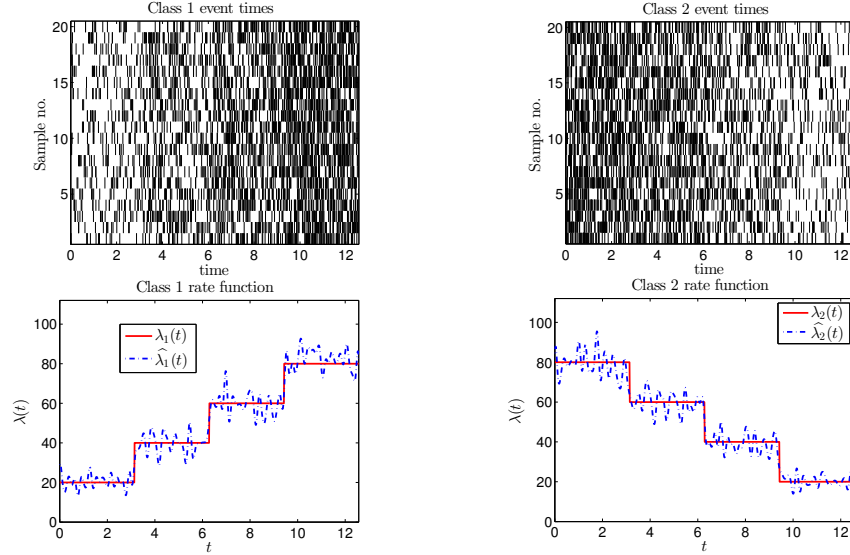


Figure 5: Event times and rate function estimates for synthetic data set 2.

Synthetic data set 3 - clustering task

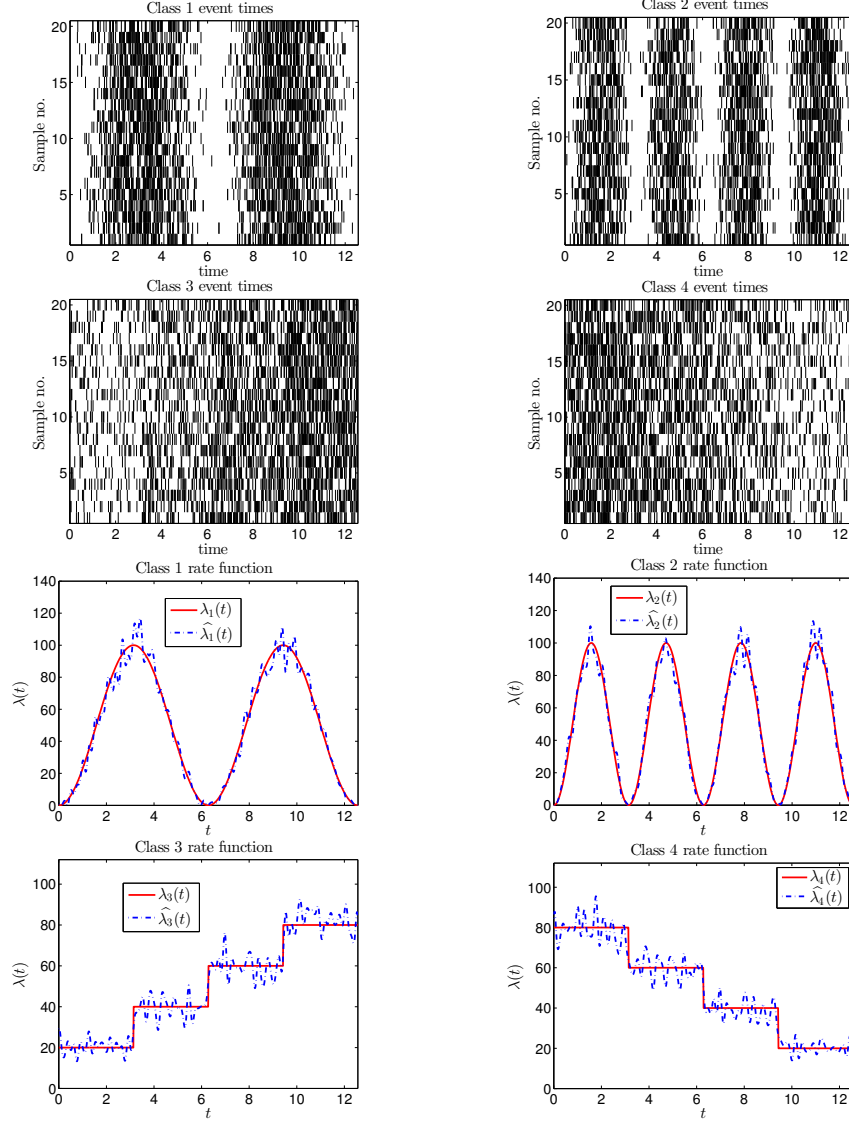


Figure 6: Results for the final synthetic data set illustrating the efficacy in our clustering method in uncovering the rate functions of the four classes.

5.2 Retail transaction data results

Next we introduce our first real data set which is made up of the till transaction times from the 6th of September 2011 to the 28th of February 2015 for 74 UK stores (10,000 transactions per store) belonging to the large UK retailer. Half of these stores have been categorised by the retailer as ‘High street’ (HS, large stores found in city centres) and the remainder as ‘travel’ (situated in airports and railway stations). These categorisations serve as the ground truth for our analysis.

5.2.1 Classification

We use repeated and stratified 5 fold cross validation (CV) (Efron and Gong, 1983) to obtain an unbiased estimate for the expected classification accuracy, true positive and true negative rate of the classifier. Specifically, samples are randomly partitioned in to 5 folds, each containing 14 or 15 samples. We ensure that the proportions of the two classes in each folder are equal, or for folds with an odd number of samples, are as equal as possible. Data in a single fold is retained as the validation set and the remainder makes up the training set which is used to obtain the NHPP estimates. This procedure is repeated 5 times with the samples in each of the folds being used once as the validation set. Finally, this process is repeated 10 times with different random partitions to give a total of 50 numerical experiments. Using high street store as the target class the average classification accuracy, true positive and true negative rates are 0.8475, 0.9943 and 0.7007 respectively. The estimates for the rate functions of each class for one CV run are shown in Figure 7. Both rate functions are periodic with the high street rate function peaking around Christmas which corresponds to the busiest shopping period on the high street. In contrast, the rate function for the travel stores (which are located in airports and railway stations) peaks around summer which coincides with the peak holiday period when holiday makers are travelling via airports and railway stations. For our results, the true positive rate is higher than the true negative rate suggesting the classifier has more of a tendency to misclassify travel stores than high street stores. Further analysis reveals that the misclassified travel stores are situated in central railways stations in large UK cities including London, Glasgow, Leeds and Liverpool. A possible explanation for the misclassification of these stores is that, although classed as travel stores, the store location means that these stores predominantly serve high street shoppers and thus have a demand profile which is similar to high street stores.

5.2.2 Clustering

Setting the number of models in the mixture to 2, and assigning each sample to the model for which its membership probability is maximal, the average clustering accuracy, true positive and true negative rates are 0.764, 1 and 0.528 respectively. The true positive rate (*resp.* true negative rate) of the clustering results is higher (lower) than the corresponding classification results. The reason for this is that the clustering procedure allows stores to naturally cluster into groups with similar temporal transaction patterns whilst for classification group membership is prescribed *a priori* in the training set. Further analysis reveals that most of the stores in the ‘travel’ cluster are located in airports. Stores categorised by the retailer as travel stores, but with temporal transaction characteristics more similar to high street stores (such as the stores situated in central railway stations in large UK cities identified in the previous section) cluster into the ‘high street’ group.

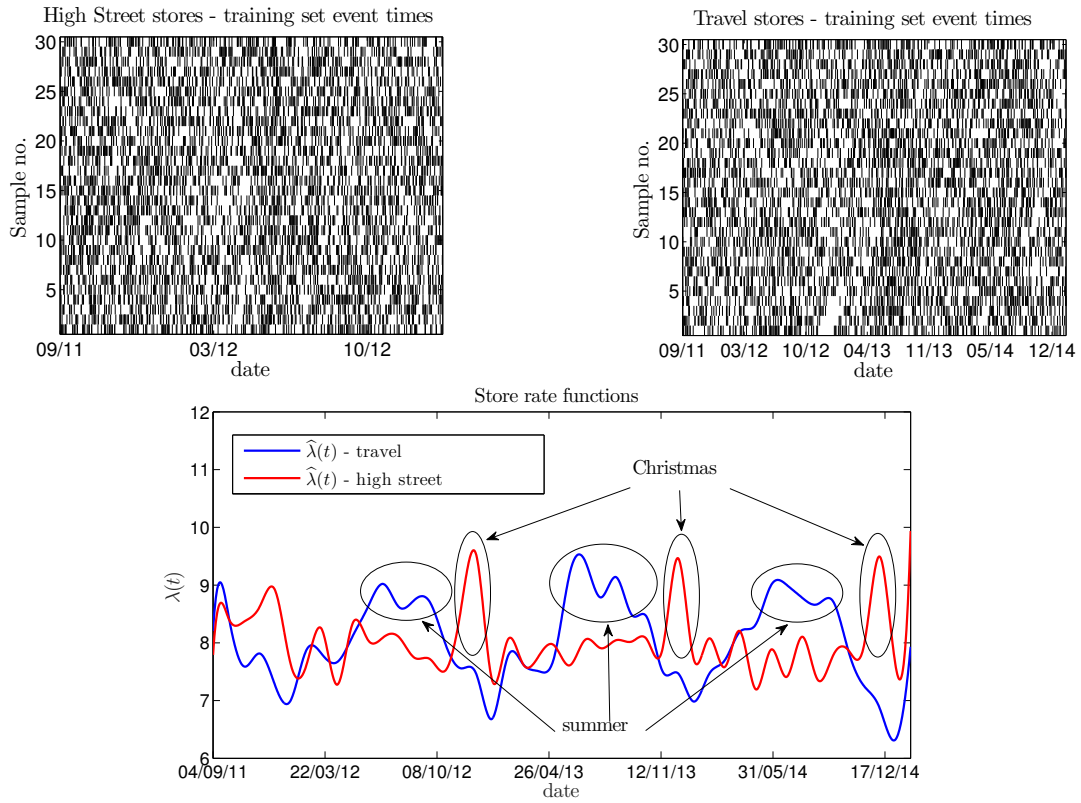


Figure 7: Raster plots of event times for a selection of samples for each class together with estimated rate functions for travel (blue) and high street (red) stores showing peak demand for travel stores is around summer while demand peaks around Christmas for high streets stores. 50 basis functions were used for the rate function estimates for each class.

The rate functions of the two NHPP models are shown in Figure 8 and clearly the blue rate function peaks around summer and corresponds to travel stores. While the red rate function peaks around Christmas and corresponds to high street stores.




Figure 8: Rate functions for the two NHPP models obtained via the clustering procedure outlined in Section 4. The blue rate function peaks around summer and the membership probabilities for most travel stores are greatest for this model. The red function peaks around Christmas and this rate function corresponds to high street stores.

Rate function estimates obtained by setting $k = 3$ are shown in Figure 9. Two clusters, with rate functions labelled 'airport 1' and 'airport 2' in the plot are comprised almost exclusively of airport stores and have very similar rate functions peaking around the summer. A third cluster, labelled 'high street and railway' is made up mainly of high street stores and railway stores located in stations in city centres and its rate function peaks around Christmas. The fact that setting $k = 3$ does not lead to more meaningful store partitions, compared to setting $k = 2$, supports the notation that the stores fall naturally in to one of two clusters. The first made up predominantly of airport stores and the second of stores located centrally in cities either on the high street or in railway stations.




Figure 9: Rate functions corresponding to the clusters for $k = 3$.

5.3 A&E data set

This data set contains the time of day (00:00:00 to 23:59:59) in every month from April 2011 to the end of December 2014 at which patients arrived at Accident and Emergency (A&E) departments in hospitals throughout England. Each patient was diagnosed as either suffering from poisoning (including alcohol poisoning) or a cardiac condition and these are used as the ground truths for the class labels. For each label there are 45 samples corresponding to each of the 45 months spanning April 2011 to the end of December 2014. Each monthly sample contains 10,000 A&E admission events.

5.3.1 Classification

Using the same CV procedure as outlined in Section 5.2.1, we obtained an average classification, true positive rate and true negative rate of 1. The estimated rate functions are shown in Figure 10 from which it can be seen than admissions for poisoning peak around midnight. This most likely due to instances of alcohol poisoning amongst revellers in the late evenings. Cardiac conditions peak in the morning which is consistent with evidence in the medical literature.(Peters et al., 1989; Muller et al., 1989)

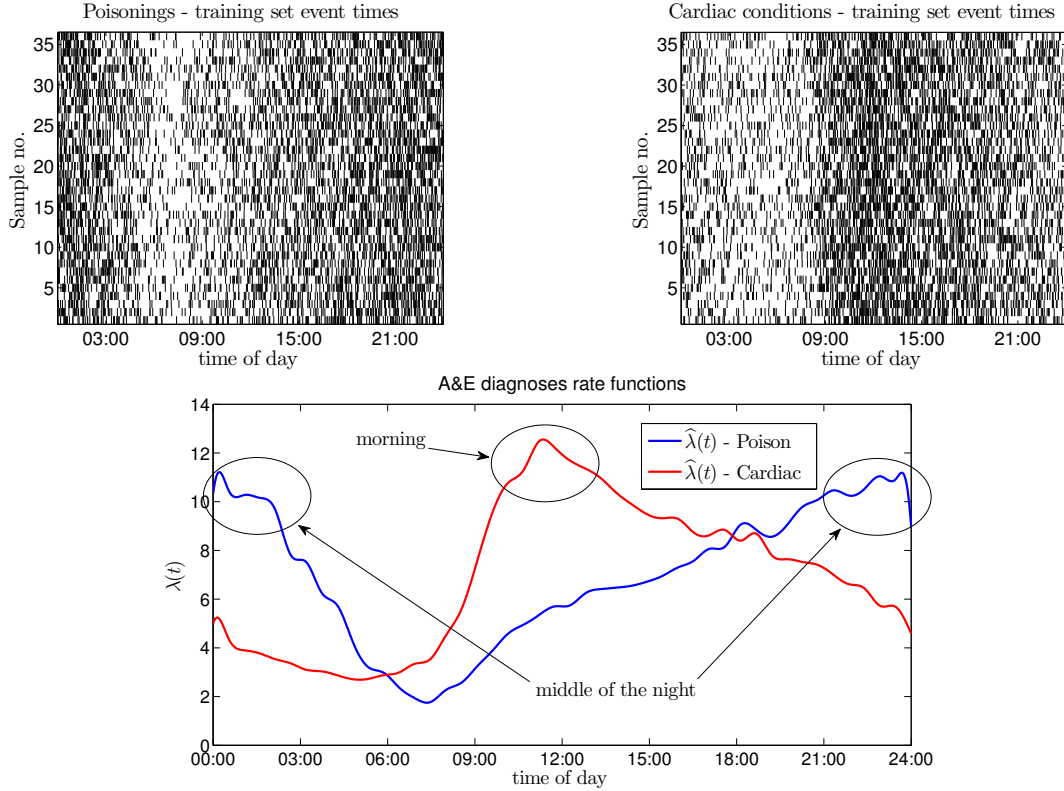


Figure 10: Event times for a selection of samples as well as estimated rate functions for poisoning (blue) and cardiac (red) diagnoses showing that poisoning instances are at a peak around midnight and cardiac conditions are maximal in the morning.

5.3.2 Clustering

Estimated rate functions can be seen in Figure 11 and these are very similar to the estimates obtained from the classification procedure with instances of poisoning peaking around night time and cardiac conditions reaching a high point in the morning. Assigning poisoning and cardiac months to the model for which their membership probabilities are maximal the clustering accuracy is 100%.




Figure 11: Estimated rate function resulting from the clustering procedure for the A&E data.

We have made Matlab code for the classification and clustering procedures outlined in Sections 3-4 and for the synthetic data set results in Section 5.1 available at <https://github.com/duncan-barrack/NHPP>. The store data and NHS A&E data are not publicly available.

6 Discussion

In this paper we have detailed principled approaches for the classification and clustering of event data using NHPP models. Results on synthetic and real data were presented which show the effectiveness of our methods. The focus of this work has been on temporal point process data observed over a fixed interval from t_1 to t_2 where $t_2 \geq t_1$. However, it could be argued that it is more natural to observe some temporal data (e.g. the times within a 24 hour period at which patients are admitted to hospital, or at which emails are sent) on a circular scale (Fisher, 1995; Jammalamadaka and Sengupta, 2001) where the notion that one temporal value is larger than another is nonsensical. Further work could investigate the applicability of our methods to these cases where circular NHPP models would need to be formed. Furthermore, our method is not restricted to temporal data and, in principle, can be generalised to multi-dimensional spatial and spatial-temporal point processes. This would provide another interesting avenue of further research.

For the clustering procedure we did not consider the problem of choosing an appropriate number of models in the NHPP mixture. In principle, the Bayesian information criterion (Schwarz et al., 1978) or Akaike information criterion (Akaike, 1974) could be used for this. Future work could investigate the suitability of these methods for choosing the number of NHPP components for our clustering procedure.

In this work, we considered samples consisting of many event times. For example, each store had 10,000 till transaction event times associated with it. For such samples obtaining an accurate estimate for the rate function is relatively straightforward as we have demonstrated in this work. However, when samples consist of a small number of event times obtaining an accurate estimate becomes more challenging and it is not clear how our procedure would perform in these circumstances. If we were to cluster individual customers according to their till transaction times for example then this is a pertinent consideration as many customers shop infrequently. Recently a new method has been proposed for rate function estimation for rare events which uses information from across many samples rather than just a single sample (Wu et al., 2013). It may be possible to integrate this method in to our clustering and classification procedures to group samples with few events.

7 Acknowledgements

We would like to thank The UK based retailer and NHS England for their support and for providing the data used in our analysis. In addition we would like to thank James Goulding for reviewing an early draft of this paper. This work was funded by EPSRC grant EP/G065802/1 - Horizon: Digital Economy Hub at the University of Nottingham and EPSRC grant EP/L021080/1 - Neo-demographics: Opening Developing World Markets by Using Personal Data and Collaboration.

References

- H. Akaike. A new look at the statistical model identification. *IEEE Trans. Autom. Control*, 19(6):716–723, 1974.
- F. Alizadeh, J. Eckstein, N. Noyan, and G. Rudolf. Arrival rate approximation by nonnegative cubic splines. *Oper Res*, 56(1):140–156, 2008.
- M.F. Arlitt and C.L. Williamson. Internet web servers: Workload characterization and performance implications. *IEEE/ACM Trans. Netw*, 5(5):631–645, 1997.
- C. Biernacki, G. Celeux, and G. Govaert. Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models. *Comput. Stat. Data Anal.*, 41(3):561–575, 2003.
- E.N. Brown, R. Barbieri, V. Ventura, R.E. Kass, and L.M. Frank. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput.*, 14(2):325–346, 2002.
- R.H. Byrd, M.E. Hribar, and J. Nocedal. An interior point algorithm for large-scale nonlinear programming. *SIAM J. Optim*, 9(4):877–900, 1999.
- R.H. Byrd, J.C. Gilbert, and J. Nocedal. A trust region method based on interior point techniques for nonlinear programming. *Math. Prog.*, 89(1):149–185, 2000.
- C. De Boor. *A practical guide to splines*, volume 27. Springer-Verlag New York, 1978.
- A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Series B Stat Methodol*, pages 1–38, 1977.
- B. Efron and G. Gong. A leisurely look at the bootstrap, the jackknife, and cross-validation. *Am. Stat.*, 37(1):36–48, 1983.
- R.F. Engle. The econometrics of ultra-high-frequency data. *Econometrica*, 68(1):1–22, 2000.
- N.I. Fisher. *Statistical analysis of circular data*. Cambridge University Press, 1995.
- T.J. Hastie, R.J. Tibshirani, and J.H. Friedman. *The elements of statistical learning : data mining, inference, and prediction*. Springer series in statistics. Springer, New York, 2009.
- J.B. Illian, E. Benson, J. Crawford, and H.J. Staines. Multivariate methods for spatial point process—a simulation study. pages 125–130. 2004.
- J.B. Illian, E. Benson, J. Crawford, and H. Staines. Principal component analysis for spatial point processes—assessing the appropriateness of the approach in an ecological context. In *Lecture Notes in Statistics 185. Case studies in spatial point process modeling*, pages 135–150. Springer, 2006.
- S.R. Jammalamadaka and A. Sengupta. *Topics in circular statistics*, volume 5. World Scientific, 2001.

- M.E. Kuhl and P.S. Bhairgond. New frontiers in input modeling: nonparametric estimation of Nonhomogeneous Poisson processes using wavelets. In *Proceedings of the 32nd conference on Winter simulation*, pages 562–571. Society for Computer Simulation International, 2000.
- S. Lee, J.R. Wilson, and M.M. Crawford. Modeling and simulation of a nonhomogeneous Poisson process having cyclic behavior. *Commun Stat Simulat*, 20(2-3):777–809, 1991.
- Peter A Lewis and Gerald S Shedler. Simulation of nonhomogeneous poisson processes by thinning. *Nav Res Log*, 26(3): 403–413, 1979.
- L.M. Linnett, D.R. Carmichael, and S.J. Clarke. Texture classification using a spatial-point process model. *IEE P-Vis Image Sign*, 142(1):1–6, 1995.
- R.D. Malmgren, D.B. Stouffer, A.E. Motter, and L.A.N. Amaral. A Poissonian explanation for heavy tails in e-mail communication. *Proc. Natl. Acad. Sci. U.S.A.*, 105(47):18153–18158, 2008.
- W.A. Massey, G.A. Parker, and W. Whitt. Estimating the parameters of a nonhomogeneous Poisson process with linear rate. *Telecommun Syst*, 5(2):361–388, 1996.
- A. McGregor, M. Hall, P. Lorier, and J. Brunskill. Flow clustering using machine learning techniques. In *Lect Notes Comput Sc*, pages 205–214. Springer, 2004.
- A.W. Moore and D. Zuev. Internet traffic classification using bayesian analysis techniques. In *ACM SIGMETRICS*, volume 33, pages 50–60. ACM, 2005.
- James E Muller, GH Tofler, and PH Stone. Circadian variation and triggers of onset of acute cardiovascular disease. *Circulation*, 79(4):733–743, 1989.
- Vern Paxson and Sally Floyd. Wide area traffic: the failure of Poisson modeling. *IEEE/ACM Trans. Netw.*, 3(3):226–244, 1995.
- R.W. Peters, J.E. Muller, S. Goldstein, R. Byington, L.M. Friedman, BHAT Study Group, et al. Propranolol and the morning increase in the frequency of sudden cardiac death (BHAT Study). *Am. J. Cardiol.*, 63(20):1518–1520, 1989.
- S.M. Ross et al. *Stochastic processes*, volume 2. John Wiley & Sons New York, 1996.
- M. Roughan, S. Sen, O. Spatscheck, and N. Duffield. Class-of-service mapping for QoS: a statistical signature-based approach to IP traffic classification. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 135–148. ACM, 2004.
- D.C. Schmittlein, D.G. Morrison, and R. Colombo. Counting Your Customers: Who-Are They and What Will They Do Next? *Manag. Sci.*, 33(1):1–24, 1987.
- G. Schwarz et al. Estimating the dimension of a model. *Ann. Stat.*, 6(2):461–464, 1978.

- David W Scott, Richard A Tapia, and James R Thompson. Nonparametric probability density estimation by discrete maximum penalized-likelihood criteria. *The annals of statistics*, pages 820–832, 1980.
- W. Thompson. *Point process models with applications to safety and reliability*. Springer Science & Business Media, 2012.
- R.A. Waltz, J.L. Morales, J. Nocedal, and D. Orban. An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Math. Prog.*, 107(3):391–408, 2006.
- S. Wu, H.G. Müller, and Z. Zhang. Functional data analysis for point processes with rare events. *Stat Sinica*, 23:1–23, 2013.
- M. Zhao and M. Xie. On maximum likelihood estimation for a general non-homogeneous Poisson process. *Scand J Stat*, pages 597–607, 1996.